

Nu@ge: Towards a solidary and responsible cloud computing service

Daniel Balouek-Thomert^{*†}, Eddy Caron^{*}, Pascal Gallard[‡], Laurent Lefevre^{*}

^{*}LIP Laboratory, UMR CNRS - ENS de Lyon - INRIA - UCB Lyon 5668

University of Lyon, France

[†]NewGeneration-SR

Paris, France

[‡]NSS Non Stop Systems

Croissy Beaubourg, France

{daniel.balouek-thomert, eddy.caron, laurent.lefevre}@ens-lyon.fr, pascal.gallard@nss.fr

Abstract—The adoption of cloud computing is still limited by several legal concerns from companies. One of those reasons is the data sovereignty, as data can be physically host in sensible locations, resulting in a lack of control for companies.

In this paper, we present the Nu@ge project aimed at building a federation of container-sized datacenter on the French territory. Nu@ge provides a software stack that enables companies to put independent datacenters in cooperation in a national mesh. Additionally, a prototype of a container-sized datacenter has been validated and patented.

Keywords: Cloud federation, Iaas, Energy-aware scheduling

I. INTRODUCTION

Cloud Computing represents a significant evolution of Information and communications technology in terms of usage and organization. It allows companies to increase their competitiveness, by lowering IT infrastructure costs and improve quality of service. This field will lead to new markets and service, remotely and on-demand accessible, led by a market growth curve of 25% per year.

Despite unquestionable benefits to users, cloud computing raises several concerns about data sovereignty. Cloud providers reveal few information about geographical location and process of data and applications. As information which has been converted and stored in binary digital form is subject to the laws of the country in which it is located, it raises several concerns from a legal standpoint. Third-party entities or governments could take control of sensible data, and legal protections may not apply if one's data is located outside of its country. Additionally, data from a specific company could be physically hosted along with data from other companies. This causes security risks in terms of attacks (attempts to access data of other companies) or marketing purposes as a secondary use.

The Nu@ge consortium¹ is composed of the following members:

- NON STOP Systems, secure cloud solutions provider and leading architect of the project

- CELESTE, Internet provider and manufacturer of the StarDC
- Oodrive, online storage specialist
- DotRiver, virtual desktop and environment provider
- Init Sys, private network operator
- New GenerationSR, Green IT consulting
- LIP6, laboratory and its research team REGAL and PHARE of UPMC University, Paris.

Those 6 innovating SMEs (Small and Medium-sized Enterprise) and 2 research teams are gathered with the ambition of creating the cloud of tomorrow : open to heterogeneous hardware and software stacks, spread on a regional network among France using low-energy consumption and ecological datacenters.

This Nu@ge project tackles several challenges. In this work, we describe a federated architecture which offers the provisioning of virtual clusters of resources over the network while providing administrators with a control over the localization of data and the Quality of Service. Our solution is based on innovative container-sized datacenters that enables deployment of a cost-effective and high-performance datacenter environment in any location, favoring the meshing of regional company-owned datacenters.

This paper is structured as follows. Section II presents the Nu@ge architecture and the associated technical choices. Section III introduces the related work. The prototype is described in Section IV. Section V discusses an evaluation in terms of energy efficiency. Section VI concludes the paper.

II. THE NU@GE APPROACH

Nu@ge defines a software stack as a coherent set of tools to homogenize the management and the exploitation of the resources. This section describes the architecture and its main components.

A. Overview

The architecture of the Nu@ge project addresses several concerns for the administrator:

- Single and shared vision of the whole infrastructure
- Simplification of services implementation

¹Nu@ge is a research project funded by the FSN (Fund for the Digital Society, BPI France) as part of the *Investissements d'Avenir* program.

- Management of a virtual cluster and an associated QoS

A strong basis of Nu@ge is the virtualization of any service. This choice implies to break the link between the logical resources and the physical resources. In particular, we consider only the QoS of virtual/logical resource without consideration of the underlying hardware. Nu@ge aims to be modular and favor the autonomy of each component. In this context, a virtual resource is characterized by its ability to be migrated depending on the following circumstances:

- Hardware defection
- Performance optimization
- Energy efficiency improvement
- Respect of QoS constraint

The unit of administration within Nu@ge is a rack. A rack contains:

- Equipments dedicated to virtualization, called V-nodes
- Equipments dedicated to storage, called IO-nodes
- Network equipments dedicated to internal communications within the rack
- Network equipments dedicated to communication with other datacenters
- Electrical equipments allowing the supervision and interventions

The high-level components and their features are described below.

B. V-node

A V-node is a physical node dedicated to the execution of virtual systems. Several services related to the infrastructures are required, among:

- Interconnection between Nu@ge and the various IaaS
- Setting up of network services
- Piloting process of electrical alimentations

In the context of Nu@ge, the main virtual machines deployed are:

The Internet Gateway Provides Internet access to the nodes, physical or virtual, present in the Nu@ge infrastructure. This specific machine enables the creation of filtering rules (firewall) in order to set a first level of security for the network services.

The VPN Gateway Offers, through the network, a secure access to Nu@ge's internal resources. Identification, authentication and data encryption are performed with digital certificates. Those are created and managed individually for each Nu@ge user.

The IaaS Gateway This is the component that links Nu@ge to the IaaS platform for the end-user. This virtual equipment is the separation between Nu@ge's area of responsibility and the IaaS administrators.

The DNS Service DNS is a primary service of Internet enabling the translation of simple identifier (ie : `http://www.nuage-france.fr`) into IP addresses (ie : 195.68.89.215). As a Nu@ge infrastructure service, it allows Internet browsing.

The Storage Access Service Creates dynamically storage units for the IaaS platforms. The storage units are available as file systems or hard drives. This service is linked to an IaaS exposing a dedicated storage zone to the Nu@ge infrastructure.

C. IO-node

The main objectives of the distributed storage system are:

- Availability
- Traceability
- Integrity and Safety

For each IaaS hosted in the Nu@ge architecture, a storage cluster is created. The number and the location of hosts depends on the contract established with the IaaS owner.

IO nodes are machines with significant storage resources. High performance disks allow improvement of current writing/reading operations while traditional disks offer bigger capacity with drawbacks of increased access time and latency.

IO nodes are connected using a dedicated subnetwork, as they need to securely exchange user's data. For that purpose, the nodes had two Gbits/Ethernet interface and an InfiniBand interface. The QoS is guaranteed, in particular during the replication of data, to ensure the resiliency in case a datacenter is suddenly not available. Additionally, the system keeps a journal of the modifications of the data.

Unlike V-nodes, an IO-node provides locally to the close V-nodes the needed storage resources. An IO-node has a high storage capacity and embed sets of hard drives. Each set can contain dozens of hard drive.

D. Communications

We use two kinds of networks within the Nu@ge architecture: the internal, dedicated to the communication between the different IaaS and, the external, used for the interconnection with the end-users.

Internal network allows the creation of private networks between the user's nodes. The setting of private networks requires an IP addressing intra and inter-datacenter, in which the flows of information are encapsulated. As the interconnection with end-user is performed via third-party internet providers, it is necessary to have several networks, depending on the segmentation set by the internet providers.

1) *External communication between the datacenters:* A simple method would consist of a star network topology, built around a central site with a full redundancy among the links. In a star topology, every node is connected to a central node called a switch. The switch acts as a server and the peripherals as the clients [1]. However, for obvious reasons of cost and architecture consistency, we do not consider that solution.

Ensuring the same continuity of service, without setting a star network topology, requires to have a number of links superior to the number of Nu@ge datacenters. Without any protocol, the interconnection of those links would cause a loop and prevents the delivery of packets.

STP (Spanning Tree Protocol) is a level 2 protocol (Ethernet) allowing the construction of an Ethernet network without

loop². STP presents a simple approach of the problem by cutting some links, to obtain a tree architecture. Due to its simple functioning, STP is widely used despite a few limitations such as the poor repartition of flows and a convergence time up to 30 seconds.

While several extensions to STP address those limitations, a new protocol named TRILL is gaining popularity. TRILL (Transparent Interconnection of Lots of Links) is a IETF standard³. This protocol presents the advantages of the routers and the network bridges by creating a level 2 network on the different links available.

Then, the protocol sets dynamic routing tables with MAC addresses. Using this level 2 routing, the protocol ensure to always have the shortest path to route the packets. In the context of Nu@ge, we use TRILL in order to manage Ethernet segmentation.

2) *Virtual Machine Mobility*: In a context of user mobility and network virtualization, getting a proper identification of an end-user over the network can be a difficult task due to the various possibility of Internet access. The protocol LISP tackles this problem by enabling migration over network while maintaining the same IP address. LISP (Locator/ID Separation Protocol) is a protocol where IP addresses have two roles: localization and identification⁴. The purpose of LISP is to solve problems related to the growing size of IPv4 routing tables. Additionally, the protocol enables users to break the link with a single internet access provider (mobile users). LISP addresses this issue by separating the localization from the identification. An IP address is used in two ways:

- Identify a machine present in a network
- Localize the identifier of the machine to route the traffic in an IP network

A distributed table of matches allows to find a locator, RLOC (Routing LOCator) from an identifier EID (Endpoint Identifier). LISP is independent of the IP addresses version and can be deployed in an incremental fashion, without the necessity of having the full Internet architecture supporting it.

E. Energy-aware management

The purpose of the energy-aware management is to evaluate the benefit of green scheduling for reducing electric consumption while matching performance objectives for the virtual machines.

The performance criteria is CPU oriented, and based on a measure of the node performance using all its CPU cores. It produces a value in flops, indicating the number of floating points operations per second. Those benchmarks are based on measurements using ATLAS⁵, HPL⁶ and Open MPI⁷.

²STP is defined in IEEE 802.1d-2004

³TRILL is defined in the RFC 6325

⁴LISP is defined in the RFC 6830

⁵Automatically Tuned Linear Algebra Software.

⁶Portable Implementation of the High-Performance Linpack Benchmark for Distributed-Memory Computers.

⁷High Performance Message Passing Library.

Regarding to the consumption criteria, two approaches are possible. A static way would imply to execute a job on all nodes before starting and measured the instantaneous electric consumption corresponding to the completion time on each node. That method is not significant for long periods because the power consumption of the machine may vary depending on the actual load or external conditions, such as the physical location of the server.

In this context, we use a more dynamic approach. The electric consumption metric is based on the number of request handled by a computational node weighted by the power consumption measured during the execution of the server. Thus, every time a client is submitting a request, each computational node will retrieve its electric consumption and total number of requests.

We rely on DIET [2], an open-source middleware that enables the execution of applications using tasks that are scheduled on distributed resources using a hierarchy of agents for scalability. DIET comprises several elements, including:

- **Client** application that uses the DIET infrastructure for remote problem solving.
- **Server Daemon (SED)** which acts as a service provider exposing functionality through a standardized computational service interface. A single SED can offer any number of computational services.
- **Agents**, deployed alone or in a hierarchy, it facilitates service location and invocation interactions between clients and SEDs. Collectively, a hierarchy of agents provides high-level and scalable services such as scheduling and data management. The head of a hierarchy is termed as **Master Agent (MA)** whereas the others are **Local Agents (LA)**.

The steps of the scheduling process are explained below:

- 1) *Submission of a virtual machine creation request*
A client issues a request describing a virtual machine template. If none of the datacenter is able to create new instances, a notification is returned to the client.
- 2) *Propagation of the request*
The request is propagated through a hierarchy of agents.
- 3) *Collect of estimation values*
Each agent computes and gathers its metrics, particularly performance and energy consumption. A reply containing those value is sent back to the scheduler.
- 4) *Sorting of candidates*
Once the scheduler retrieves all the replies, it proceeds to a sort according to a specific criteria. The first ranked node is then elected and notified.
- 5) *Virtual machine creation*
The virtual machine is created on the elected node.

Others criteria exists in the literature, involving the consideration of idle consumption [3] or the use rate [4] of the physical nodes.

We coupled the scheduling process to the provisioning of resources while taking into account energy-related events such as fluctuations of electricity prices or heat peaks. This feature

[5] enable autonomic decisions from the scheduler by checking pre-defined threshold before executing placement/provisioning decisions.

III. RELATED WORK

This section presents the related work in the context of cloud providers and the existing cloud federation approaches. Then, we detail the notion of modular datacenters and its expected features. Finally, we justify the choice of implementing our own distributed storage solution and the different of energy-aware scheduling mechanisms.

A. Cloud Providers

It exists plenty of surveys related to cloud providers [6] [7]. Table I shows a classification of the current major public IaaS providers. That list of providers is by no means exhaustive, but it includes those that are believed to be the current major players at least in the European and North American markets [8]. Most of the Cloud providers imposes to operate according to its specific models and protocols. This problem is known as vendor lock-in and restrict the transition and interoperability across providers [9]. Furthermore, the headquarters and datacenters location columns show that most providers are based in the USA while only a few are based in the Europe. This constitutes one of Nu@ge’s motivation which aims to relocate Cloud Computing in France.

A first step towards solving those problem is provided by IaaS stacks, such as OpenStack [9] or VMWare vCloud [10], for creating and managing infrastructure cloud services in private, public, and hybrid clouds.

B. Cloud federation approach

The Cloud federation approach [10] aims to resolve the issues of both providing a unified platform for managing resources at different levels and abstracting interaction models of different cloud providers. Several European projects are providing stacks and/or adaptation of cloud-based systems at IaaS levels. Contrail [11], [12] aims at solving the vendor lock-in problem by allowing the seamless switch of cloud provider. InterCloud [10] is a federated cloud computing environment that aims at provisioning application in a scalable computing environment, achieving QoS under variable workload, resource and network conditions. In the Reservoir project [13], the authors propose an architecture for an open federated cloud computing platform. In such architecture, each resource provider is an autonomous entity with its own business goals. Celesti et al. [14], proposed the Dynamic Cloud Collaboration, an approach for setting up highly dynamic cloud federations. In this work, a distributed agreement must be reached among the already federated partners to dynamically federate a new provider.

C. Modular Datacenter

Clouds depend on datacenters, large facilities used to house computer systems and associated components, such as telecommunications and storage systems. A modular data

center system is a portable method of deploying data center capacity. As an alternative to the traditional data center, a modular data center can be placed wherever data capacity is needed.

Modular data center systems consist of purpose-engineered modules and components to offer scalable data center capacity with multiple power and cooling options. Numerous manufacturers such as Google, IBM, Sun, Verrari or HP built modular datacenters into standard intermodal containers (shipping containers) with the following key features:

High Density Maximum accommodation of servers, storage and network equipments within a limited surface.

Cost Reduction By comparison to the building and exploitation of a traditional raised-floor data center.

Self-contained Cooling Self-contained cooling technologies, which can enable a cost savings and improve system reliability.

Environmentally Responsible Minimal carbon footprint.

Disaster Recovery and Security Characterized by the time of autonomy of the container and the physical equipments dedicated to ensure its integrity.

Fast deployment Usually expected to be less than 6 months to be put in service after order to the manufacturer.

Industry relies on the TIA-942 specification [15] to classify the minimum requirements for telecommunications infrastructure of data centers and computer rooms into 4 categories, presented in Table II.

	Characteristics of the site infrastructure design topology	Theoretical availability
Tier 1	Single path for power and cooling distribution. No redundant components.	99.671%
Tier 2	Single path for power and cooling distribution. Includes redundant components	99.741%
Tier 3	Multiple power and cooling distribution paths. Only one active path. Includes redundant components. Concurrently maintainable.	99.982%
Tier 4	Multiple power and cooling distribution paths. All paths are active. Includes redundant components. Concurrently maintainable. Fault tolerant.	99.995%

TABLE II
TIA-942 TIER SYSTEM

D. Distributed storage

The Nu@ge project specifications insists on resiliency. In case of the loss of connectivity of a datacenter, the data storage must be distributed among the federation while traceability, integrity and security of data must be ensured. Additionally, the storage system must keep a journal of data modifications to retrieve a coherent state after a incident. The following part evaluates existing distributed storage solutions with the purpose of integrating one suiting Nu@ge needs. Apart from

Provider	Headquarters	Datacenters location
Amazon AWS	USA	USA, Brazil, Ireland, Japan, Singapore, Australia
AT&T Cloud	USA	USA
Google Compute Engine	USA	USA, UK
Hosting.com	USA	USA
GoGrid	USA	USA
Microsoft Windows Azure	USA	USA, Ireland, Netherlands, Hong Kong, Singapore
Rackspace	USA	USA, UK, Hong Kong
OpSource	USA	USA, France, UK
Terramark	USA	USA, Canada, Brazil, Colombia, Dominican Republic, Belgium, France, Germany, Ireland, Italy, Luxembourg, Netherlands, Spain, Sweden, Turkey, UK, China, Japan, Singapore, Australia
Softlayer	USA	USA, Netherlands, Singapore
Aruba Cloud	Italy	Italy
CloudSigma	Switzerland	Switzerland, USA
Gandi	France	France, USA
GreenQloud	Iceland	Iceland
Lunacloud	UK	France, Germany, Latvia, Portugal
CloudWatt	France	France
Numergy	France	France

TABLE I
CLOUD PROVIDERS.

functional aspects, the solution must be mature enough to be considered for production uses.

There is two main categories of storage [16] [17], Network Attached Storage (NAS) and Storage Area Network (SAN). Network Attached Storage is file-level computer data storage server connected to a computer network providing data access to a heterogeneous group of clients, while Storage Area Network is a dedicated network that provides access to consolidated, block level data storage.

1) *Network attached storage*: The Ceph project [18] is quite close from our requirements. However, a Ceph cluster can only handle one file system, which is a serious technical restriction.

HDFS [19] is conceived to distribute computations between several nodes. One of the nodes, the *namenode* is a necessary gateway to the system. It constitutes a serious bottleneck and is inappropriate for Nu@ge architecture.

GlusterFS, MooseFS, Pohmelfs and XtremFS presented various limitations. Unstability reasons, troubles with operating system support or lack of contribution support led us to exclude those projects from our choices.

2) *Storage area network*: Despite its lack of journalization, Ceph project [18] features an extensive block data storage. Nevertheless, Ceph cluster gives no information about the localization of data. In this context, data traceability, one of the main objective of Nu@ge, could be achieved only by creating a Ceph cluster per datacenter. This solution is not worth considering due to the high resource consumption of Ceph. Sheepdog [20] seems relatively inactive and only works with QEMU/KVM virtualization technologies. Some of Sheepdog technical choices would lead to scalability problems in terms of storage or number of datacenters.

Unfortunately, there were no project fitting our requirements among those two categories. In particular, the notion of journal is not a common feature and would be hard to implement

among an existing file system. We decide to initiate a new project over a SAN, as it is less complex to implement. Unlike a NAS that needs the installation of a software on the client desktop, block level data storage can be access through standard protocols (specifically iSCSI [21], supported in a native fashion by numerous operating systems).

E. Energy-aware scheduling

Despite the increasing popularity of Cloud computing, infrastructures on which they rely are seldom fully utilized [22], mostly as result of overprovisioning to handle peak demands. Workloads with large variations in demand can lead to periods of low resource utilization. As resources are generally not energy proportional, meaning their power consumption at low load is already high, the energy efficiency of an infrastructure is reduced during such periods. Power saving techniques proposed to circumvent such problems consist in slowing down certain server components [23], [24] during periods of light load — techniques that according to Le Sueur *et al.* are becoming less attractive on modern hardware [25] — or using software techniques to put idle servers into low power consumption modes [26], [27]. These techniques are well suited to Clouds where virtualization is mainstream.

In the context of Nu@ge, the authors provide techniques for assigning virtual machines to federated resources, by exploring energy efficient resource provisioning. We defined mechanisms to adapt resource allocation according to energy-related events and administrator-defined rules [5]. We take benefit of the plug-in scheduler functionality of the DIET middleware [2] to evaluate.

IV. PROTOTYPE

This section briefly describes a prototype implementation of the Nu@ge project, that follows the architecture described in Figure 2. Such implementation has been used for evaluating performance and feasibility of the proposed approach. The



Fig. 1. The public presentation of the StarDC container occurred on September 18th 2014 during Nu@ge inauguration in CELESTE headquarters, Marne-la-vallee, France.

prototype has been realized by using a customized version of OpenStack over 4 different geographic locations in France.

A. Roles

The IaaS administrator is in charge of the virtual infrastructure offered by Nu@ge. The Nu@ge administrator sets and configure resources (physical or virtual nodes, storage disks, virtual routers...).

B. StarDC

The StarDC (Figure 1) features 4 service units of 19 inches racks and can hold up to 168 computing servers. The container facility spans 15 square meters, a power capacity of 18 kilowatts and a Power Usage Effectiveness (PUE) of 1.24. The *StarDC* is built within Tier 3 specifications and is the subject of a patent.

Unlike most of the modular datacenters, the StarDC does not use water cooling. It allows an easier implantation, a broader range of physical locations and an eco-responsible behavior since ambient air is used to cool the container. The StarDC uses a mechanism of temperature uphold using outdoor air as a free cooling source. The purpose is to take advantage of outdoor temperatures to obtain a natural chilling of equipments. When the air is injected into machines, its temperature raises by a delta number of 10° (common value among commercialized servers). In case the outdoor temperature is higher than required, we use air conditioning to chill it.

The Nu@ge customer is in charge of the setting of the cold aisle temperature. If he chooses a temperature of 20° to favor confort of use, the air conditioning will be active approximately 20% of the year (varies depending on the location). Choosing a temperature value up to 25° and more results in less air conditioning and a better ecological impact.

C. Construction of a IaaS

The creation of a new IaaS does not impact the architecture of Nu@ge. The main evolution concerns the evolution of the virtual nodes allowing the mutualisation of physical resources. In particular, the instantiation of:

- A storage access point

- An IaaS access point
- A virtual switch interconnecting the IaaS equipments

As a result, several IO-nodes and V-nodes can be used by multiple IaaS.




Fig. 2. Nu@ge architecture including gateways and a IaaS

D. Storage cluster

Nu@ge racks contains two IO-nodes. As storage management can require large computational resources, an IO-node features dual-core CPUs for a total of 24 threads and 256 gigabytes of RAM. Deployment of IO-nodes is performed via the following steps:

- 1) Booting via PXE / TFTP protocols
- 2) Configuration using Puppet
- 3) Creation and configuration of an object storage in RAID1
- 4) Creation and configuration of RAID6 objects
- 5) Creation of logical volumes

At this point, the node executes an Openstack storage service specific to the newly created IaaS, and the storage server. This organization is coherent with Nu@ge objective of data isolation between IaaS and data traceability for the administrators.

V. DISCUSSION

The Power Usage Effectiveness is a usual metric to evaluate the green quality of a datacenter [28]. From a practical point of view, it is a measure of how efficiently a computer data center uses energy; specifically, how much energy is used by

the computing equipment (in contrast to cooling and other overhead). The PUE is expressed by the ratio:

$$PUE = \frac{TotalFacilityPower}{ITEquipmentPower} \quad (1)$$

Nevertheless, it is very hard to know the real PUE from a company because the area of Equipment Power can be debatable. As an example, for the Google Data center, if Google considers only the servers and the air conditioning, Google gives a PUE of 1.06. But if Google add generators, transformers, site substations and natural gas then the PUE is 1.14.

Green Datacenter from green.ch company (Switzerland) was designed with energy efficiency and reduction consideration. This project is based on energy-optimized data center architecture, latest generation of air conditioners, heat exchangers, waste heat utilization in new office building.

The container-sized datacenter designed by Nu@ge aims at keeping the PUE under the value of 1.30, using two cooling operating modes:

- Total free cooling when the room temperature is in-range with the servers specifications. That range is set by the customer resulting in a PUE value of 1,16.
- Air recycling with air conditioning when the temperature is out of range results in a PUE value of 1,55.

Thus, the PUE relies strongly on the climate conditions, and customer-defined rules. In the case of Nu@ge's *StarDC* at Marne-La-Vallee (France), weather forecast indicates that 80% of the time, the temperature is below 23°. The theoretical maximal value for the PUE is then:

$$PUE_{Nu@ge} = 80\% \times 1.16 + 20\% \times 1.55 = 1.24 \quad (2)$$

Data Center	Company	PUE
Prineville DC	Facebook	1.07
Google DC	Google	1.14
StarDC	Nu@ge	1.24
Green Datacenter	grench.ch	1.4

TABLE III

PUE COMPARISON OF DIFFERENT DATACENTERS. THOSE VALUES ARE GIVEN BY EACH PROJECT BUT NO INDEPENDENT EVALUATION WAS DONE.

VI. CONCLUSION

The Nu@ge project aims at designing and building a network of modular datacenters dedicated to the virtualization of IT services. This architecture can be referred as "Cloud" with the following guarantees:

- Control over the underlying infrastructure
- Knowledge of data location
- Control over the different QoS

By using Nu@ge, a final user (ie. the administrator of an IT system) is focus on the management of a virtual cluster seamlessly spread on a collection of datacenters with a support on the infrastructure supervision.

ACKNOWLEDGEMENTS

Nu@ge is a research project funded by the FSN (Fund for the Digital Society, BPI France) as part of the *Investissements d'Avenir* program.

The authors acknowledge the support of the Nu@ge consortium, composed of the following members: NON STOP Systems, secure cloud solutions provider; CELESTE, Internet provider and manufacturer of the StarDC; Oodrive, online storage specialist; DotRiver, private network operator; New GenerationSR, Green IT consulting; LIP6, laboratory and its research team REGAL and PHARE of UPMC University, Paris.

REFERENCES

- [1] L. G. Roberts and B. D. Wessler, "Computer network development to achieve resource sharing," in *Proceedings of the May 5-7, 1970, Spring Joint Computer Conference*, ser. AFIPS '70 (Spring). New York, NY, USA: ACM, 1970, pp. 543-549.
- [2] E. Caron and F. Desprez, "DIET: A scalable toolbox to build network enabled servers on the grid," *International Journal of High Performance Computing Applications*, vol. 20, no. 3, pp. 335-352, 2006.
- [3] M. E. M. Diouri, O. Glück, L. Lefèvre, and J.-C. Mignot, "Your cluster is not power homogeneous: Take care when designing green schedulers!" in *IGCC-4th IEEE International Green Computing Conference*, 2013.
- [4] S.-H. Lim, B. Sharma, B.-C. Tak, and C. R. Das, "A dynamic energy management scheme for multi-tier data centers," in *ISPASS*. IEEE Computer Society, 2011, pp. 257-266.
- [5] D. Balouek-Thomert, E. Caron, and L. Lefevre, "Energy-aware server provisioning by introducing middleware-level dynamic green scheduling," in *Workshop HPPAC'15. High-Performance, Power-Aware Computing*. Hyderabad, India: In conjunction with IPDPS 2015, May 2015.
- [6] B. P. Rimal, E. Choi, and I. Lumb, "A taxonomy and survey of cloud computing systems," in *INC, IMS and IDC, 2009. NCM'09. Fifth International Joint Conference on*. Ieee, 2009, pp. 44-51.
- [7] R. Prodan and S. Ostermann, "A survey and taxonomy of infrastructure as a service and web hosting cloud providers," in *Grid Computing, 2009 10th IEEE/ACM International Conference on*. IEEE, 2009, pp. 17-25.
- [8] N. Ferry, A. Rossini, F. Chauvel, B. Morin, and A. Solberg, "Towards model-driven provisioning, deployment, monitoring, and adaptation of multi-cloud systems," in *CLOUD 2013: IEEE 6th International Conference on Cloud Computing*, 2013, pp. 887-894.
- [9] B. Satzger, W. Hummer, C. Inzinger, P. Leitner, and S. Dustdar, "Winds of change: From vendor lock-in to the meta cloud," *IEEE internet computing*, vol. 17, no. 1, pp. 69-73, 2013.
- [10] R. Buyya, R. Ranjan, and R. N. Calheiros, "Intercloud: Utility-oriented federation of cloud computing environments for scaling of application services," in *Algorithms and architectures for parallel processing*. Springer, 2010, pp. 13-31.
- [11] R. G. Cascella, C. Morin, P. Harsh, and Y. Jegou, "Contrail: A reliable and trustworthy cloud platform," in *Proceedings of the 1st European Workshop on Dependable Cloud Computing*, ser. EWDC '12. New York, NY, USA: ACM, 2012, pp. 6:1-6:2.
- [12] E. Carlini, M. Coppola, P. Dazzi, L. Ricci, and G. Righetti, "Cloud federations in contrail," in *Euro-Par 2011: Parallel Processing Workshops*. Springer, 2012, pp. 159-168.
- [13] B. Rochwerger, D. Breitgand, E. Levy, A. Galis, K. Nagin, I. M. Llorente, R. Montero, Y. Wolfsthal, E. Elmroth, J. Caceres *et al.*, "The reservoir model and architecture for open federated cloud computing," *IBM Journal of Research and Development*, vol. 53, no. 4, pp. 4-1, 2009.
- [14] A. Celesti, F. Tusa, M. Villari, and A. Puliafito, "How to enhance cloud architectures to enable cross-federation," in *Cloud Computing (CLOUD), 2010 IEEE 3rd International Conference on*. IEEE, 2010, pp. 337-345.
- [15] Telecommunication Industry Association, "Tia-942 data center standards overview," *White Paper*, 2006.
- [16] D. Sacks, "Demystifying storage networking das, san, nas, nas gateways, fibre channel, and iscsi," *IBM Storage Networking*, pp. 3-11, 2001.
- [17] G. A. Gibson and R. Van Meter, "Network attached storage architecture," *Communications of the ACM*, vol. 43, no. 11, pp. 37-45, 2000.

- [18] S. A. Weil, S. A. Brandt, E. L. Miller, D. D. E. Long, and C. Maltzahn, "Ceph: A scalable, high-performance distributed file system," in *Proceedings of the 7th Symposium on Operating Systems Design and Implementation (OSDI)*, 2006, pp. 307–320.
- [19] K. Shvachko, H. Kuang, S. Radia, and R. Chansler, "The hadoop distributed file system," in *Proceedings of the 2010 IEEE 26th Symposium on Mass Storage Systems and Technologies (MSST)*, ser. MSST '10. Washington, DC, USA: IEEE Computer Society, 2010, pp. 1–10.
- [20] K. Morita, "Sheepdog: Distributed storage system for qemu/kvm," *LCA 2010 DS&R miniconf*, 2010.
- [21] J. Satran, K. Meth, C. Sapuntzakis, M. Chadalapaka, and E. Zeidner, "Internet small computer systems interface (iscsi)," United States, 2004.
- [22] A. Hawkins, "Unused servers survey results analysis," *The Green Grid*, 2010.
- [23] E. V. Carrera, E. Pinheiro, and R. Bianchini, "Conserving disk energy in network servers," in *Proceedings of the 2003 International Conference on Supercomputing (ICS-03)*. New York: ACM Press, Jun. 23–26 2003, pp. 86–97.
- [24] D. C. Snowdon, S. Ruocco, and G. Heiser, "Power management and dynamic voltage scaling: Myths and facts," in *Proceedings of the 2005 Workshop on Power Aware Real-time Computing*, Sep. 2005.
- [25] E. Le Sueur and G. Heiser, "Dynamic voltage and frequency scaling: The laws of diminishing returns," in *Proceedings of the 2010 international conference on Power aware computing and systems*. USENIX Association, 2010, pp. 1–8.
- [26] A. Beloglazov and R. Buyya, "Managing overloaded hosts for dynamic consolidation of virtual machines in cloud data centers under quality of service constraints," *IEEE Trans. Parallel Distrib. Syst.*, vol. 24, no. 7, pp. 1366–1379, 2013.
- [27] E. Feller, L. Rilling, and C. Morin, "Snooze: A scalable and autonomic virtual machine management framework for private clouds," in *CCGRID*. IEEE, 2012, pp. 482–489.
- [28] C. Belady, A. Rawson, J. Pflueger, and T. Cader, "Green grid data center power efficiency metrics: PUE and DCIE," The Green Grid, Tech. Rep. White Paper 6, 2008.